

Technical Bulletin #0295
DB2 VERSION 7 RELEASE 1

Issued Date: December 20, 2001
Effective Date: December 20, 2001
Section/Groups: Software Management/Database Administration
Submitted By: Brian Triptow
Approved By: Jim Calaway

In January, 2002, the State of Utah plans to migrate to DB2 Version 7 Release 1 on CPU4 in the DBD1 development subsystem. This software will impact ORSIS WERD, WERI, WERA, and WERT regions. Subsequent installation on CPU3 is tentatively scheduled for late March, 2002, and will impact ORSIS production WERP.

All DB2 PLANS and PACKAGES will be rebound, and PLAN_TABLES will be altered for Version 7.

New version features can be exploited after the time when fallback to Version 6 would not be a consideration.

SUMMARY OF CHANGES TO DB2 FOR OS/390 AND Z/OS VERSION 7
(formatted from IBM Installation Guide)

DB2 for OS/390 and z/OS Version 7 delivers an enhanced relational database server solution for OS/390. This release focuses on greater ease and flexibility in managing your data, better reliability, scalability, and availability, and better integration with the DB2 family.

Enhancements for Managing Data

Version 7 delivers the following enhancements for managing data:

- DB2 now collects a comprehensive statistics history that:
 - lets you track changes to the physical design of DB2 objects; and,
 - lets DB2 predict future space requirements for table spaces and indexes more accurately and runs utilities to improve performance.
- Database administrators can now manage DB2 objects more easily and no longer must maintain their utility jobs (even when new objects are added) by using enhancements that let them:
 - dynamically create object lists from a pattern-matching expression; and,
 - dynamically allocate the data sets that are required to process those objects.
- More flexible DBADM authority lets database administrators create views for other users.
- Enhancements to management of constraints let you specify a constraint at the time you create primary or unique keys. A new restriction on the DROP INDEX statement requires that you drop the primary key, unique key, or referential constraint before you drop the index that enforces a constraint.

Enhancements for Reliability, Scalability, and Availability

tb0295.doc 2001-12-14

Version 7 delivers the following enhancements for the reliability, scalability, and availability of your e-business:

- The DB2 Utilities Suite provides utilities for all of your data management tasks that are associated with the DB2 catalog.
- The new UNLOAD utility lets you unload data from a table space or an image copy data set. In most cases, the UNLOAD utility is faster than the DSNTIAUL sample program, especially when you activate partition parallelism for a large partitioned table space. UNLOAD is also easier to use than REORG UNLOAD EXTERNAL.
- The new COPYTOCOPY utility lets you make additional image copies from a primary image copy and registers those copies in the DB2 catalog. COPYTOCOPY leaves the target object in read/write access mode (UTRW), which allows Structured Query Language (SQL) statements and some utilities to run concurrently with the same target objects.
- Parallel LOAD with multiple inputs lets you easily load large amounts of data into partitioned table spaces for use in data warehouse applications or business intelligence applications. Parallel LOAD with multiple inputs runs in a single step, rather than in different jobs.
- A faster online REORG is achieved through the following enhancements:
 - Online REORG no longer renames data sets, which greatly reduces the time that data is unavailable during the SWITCH phase.
 - Additional parallel processing improves the elapsed time of the BUILD2 phase of REORG SHRLEVEL(CHANGE) or SHRLEVEL(REFERENCE).
- More concurrence with online LOAD RESUME is achieved by letting you give users read and write access to the data during LOAD processing so that you can load data concurrently with user transactions.
- More efficient processing for SQL queries:
- More transformations of subqueries into a join for some UPDATE and DELETE statements.
 - Fewer sort operations for queries that have an ORDER BY clause and WHERE clauses with predicates of the form COL=constant.
 - More parallelism for IN-list index access, which can improve performance for queries involving IN-list index access.
- The ability to change system parameters without stopping DB2 supports online transaction processing and e-business without interruption.
- Improved availability of user objects that are associated with failed or canceled transactions:
 - You can cancel a thread without performing rollback processing.
 - Some restrictions imposed by the restart function have been removed.
 - A NOBACKOUT option has been added to the CANCEL THREAD command.

- Improved availability of the DB2 subsystem when a log-read failure occurs. DB2 now provides a timely warning about failed log-read requests and the ability to retry the log read so that you can take corrective action and avoid a DB2 outage.
- Improved availability in the data sharing environment:
 - Group attachment enhancements let DB2 applications generically attach to a DB2 data sharing group.
 - A new LIGHT option of the START DB2 command lets you restart a DB2 data sharing member with a minimal storage footprint, and then terminate normally after DB2 frees the retained locks that it can.
 - You can let changes in structure size persist when you rebuild or reallocate a structure.
- Additional data sharing enhancements include:
- Notification of incomplete units of recovery.
 - Use of a new OS/390 and z/OS function to improve failure recovery of group buffer pools.
- An additional enhancement for e-business provides improved performance with preformatting for INSERT operations.

Easier Development and Integration of e-Business Applications

Version 7 provides the following enhancements, which let you more easily develop and integrate applications that access data from various DB2 operating systems and distributed environments:

- DB2 XML Extender for OS/390 and z/OS, a new member of the DB2 Extender family, lets you store, retrieve, and search XML documents in a DB2 database.
- Improved support for UNION and UNION ALL operators in a view definition, a nested table expression, or a subquery predicate, improves DB2 family compatibility and is consistent with SQL99 standards.
- More flexibility with SQL gives you greater compatibility with DB2 on other operating systems:
 - Scrollable cursors let you move forward, backward, or randomly through a result table or a result set. You can use scrollable cursors in any DB2 applications that do not use DB2 private protocol access.
 - A search condition in the WHERE clause can include a subquery in which the base object of both the subquery and the searched UPDATE or DELETE statement are the same.
 - A new SQL clause, FETCH FIRST n ROWS, improves performance of applications in a distributed environment.
 - Fast implicit close in which the DB2 server, during a distributed query, automatically closes the cursor when the application attempts to fetch beyond the last row.
 - Support for options USER and USING in a new authorization clause for CONNECT statements lets you easily port applications that are developed on the workstation to DB2 for OS/390 and z/OS. These options also let applications that run under WebSphere to reuse DB2 connections for different users and to enable DB2 for OS/390 and z/OS to check passwords.

- For positioned updates, you can specify the FOR UPDATE clause of the cursor SELECT statement without a list of columns. As a result, all updatable columns of the table or view that is identified in the first FROM clause of the fullselect are included.
- A new option of the SELECT statement, ORDER BY expression, lets you specify operators as the sort key for the result table of the SELECT statement.
- New datetime ISO functions return the day of the week with Monday as day 1 and every week with seven days.
- Enhancements to Open Database Connectivity (ODBC) provide partial ODBC 3.0 support, including many new application programming interfaces (APIs), which increase application portability and alignment with industry standards.
- Enhancements to the LOAD utility let you load the output of an SQL SELECT statement directly into a table.
- A new component called Precompiler Services lets compiler writers modify their compilers to invoke Precompiler Services and produce an SQL statement coprocessor. An SQL statement coprocessor performs the same functions as the DB2 precompiler, but it performs those functions at compile time. If your compiler has an SQL statement coprocessor, you can eliminate the precompile step in your batch program preparation jobs.
- Support for Unicode-encoded data lets you easily store multilingual data within the same table or on the same DB2 subsystem. The Unicode encoding scheme represents the code points of many different geographies and languages.

Improved Connectivity

Version 7 offers improved connectivity:

- Support for COMMIT and ROLLBACK in stored procedures lets you commit or roll back an entire unit of work, including uncommitted changes that are made from the calling application before the stored procedure call is made.
- Support for Windows Kerberos security lets you more easily manage workstation clients who seek access to data and services from heterogeneous environments.
- Global transaction support for distributed applications lets independent DB2 agents participate in a global transaction that is coordinated by an XA-compliant transaction manager on a workstation or a gateway server (Microsoft Transaction Server or Encina, for example).
- Support for a DB2 Connect Version 7 enhancement lets remote workstation clients quickly determine the amount of time that DB2 takes to process a request (the server elapsed time).
- Additional enhancements include:
 - Support for connection pooling and transaction pooling for IBM DB2 Connect.
 - Support for DB2 Call Level Interface (DB2 CLI) bookmarks on DB2 UDB for UNIX, Windows, OS/2.

Features of DB2 for OS/390 and z/OS

tb0295.doc 2001-12-14

Version 7 of DB2 UDB Server for OS/390 and z/OS offers several features that help you integrate, analyze, summarize, and share data across your enterprise:

- QMF Version 7.
- QMF for Windows.
- Net Search Extender for in-memory text search for e-business applications.